

WHAT IS A DATABASE RELIABILITY ENGINEER (DBRE)?



As today's data driven businesses continue to utilize distributed datastores spread across the cloud, the Database Reliability Engineer (DBRE) role has become increasingly important to keeping this data easily accessible and safe. The DBRE leads the planning, managing, and scaling of these datastores to ensure a business' complex data requirements are met and it can easily access its data in a fast, reliable, and safe manner.

Working with critical data can create reservations about how a business utilizes it. Often the fear of losing important information can stifle innovation, yet innovation is what keeps businesses thriving in a challenging world full of fearless startups and unforeseen economic turns. These fears develop into gatekeeping and risk adverse policies engineers must deal with in order to utilize the data they need to create and innovate.

In the [DevOps world](#), the DBRE ensures the quality of orchestration and integration of tools needed to support daily operations by patching together existing infrastructure with cloud solutions and additional data infrastructures. Failure to do so in a sound way can justify the hassles engineers face when sourcing the data they need because of the fear of creating a costly outage and crippling operations.

Making datastores openly accessible and resilient

One of the challenges DBREs face is how they make data open and easy to access while providing confidence the databases they are in charge of will work under any conditions. This involves

protecting the data through rigorous testing of backup and recovery processes and frequently auditing well-regulated security procedures. These tasks, although important and necessary, are also quite redundant. The path to alleviating fear while increasing velocity is paved in automation. DBREs use software and tooling to automate manual tasks and enable engineers to move fast without concern of losing data during their experiments.

By automating crucial operations like failovers, backups, and resource provisions, organizations can feel confident they are safe from catastrophic data losses. DBREs create and support usage guardrails by communicating general data protection regulation (GDPR) practices to developers and engineers who can then work within datastores without facing slow downs from gatekeepers or the organization's overwrought fear of damaging the database. Engineers should be assured that if their experiment goes wrong, someone with a backup who knows how to use it is only a quick high priority email away.


Standardization, Automation and Resilience

DBREs use [service level objectives \(SLOs\)](#) and risk analysis to determine which problems to address and which problems to automate. They must be familiar not only with the data they are sourcing but be able to evaluate the multiple tools available to manage an organization's data by choosing the tools and processes which will help them create resilient systems, encouraging everyone to work with the datastore as much as possible. This framework is developed into a reliability program DBREs are responsible for maintaining along with the greater devops and leadership teams to achieve high levels of standardization, reliability, testability, maintainability, and system availability.

4 Reliability Engineering Objectives for Data Protection

- Apply engineering knowledge and specialist techniques to prevent or to reduce the likelihood or frequency of failures.
- Identify and correct the causes of failures that do occur despite the efforts to prevent them.
- Determine ways of coping with failures that do occur, if their causes have not been corrected.
- Applying methods for estimating the likely reliability of new designs and for analyzing reliability data.

4 Reliability Engineering Objectives for Data Protection



While testing and developing the data protection frameworks, a key task for any DBRE is to identify and eliminate “toil”. The Google SRE teams often use the phrase “eliminate toil” and define toil as “the kind of work tied to running a production service that tends to be manual, repetitive, automatable, tactical, devoid of enduring value, and that scales linearly as a service grows.” Effectively eliminating toil is done through automation and standardization of manual work which is repetitive, non-creative, and non-challenging.

Toil is the type of work the right software can help solve by quickly capturing database configurations, obtaining visibility into out-of-compliance conditions, and taking corrective action with less manual effort. BMC’s [Bladelogic Database Automation](#) is designed to consistently replicate task execution through improving speed, agility, and release quality, while lowering risk, and without adding staff. Automating processes like these will streamline database workflows and empower developers to create and innovate. The right software can help DBREs manage compliance issues, heighten security, increase data quality, and lower costs when implemented effectively. This leaves DBREs left to anticipate future problems instead of reacting to current ones.

Breaking down silos

Databases are not meant to be silos of information separate from the rest of an organization's operational components. Managing databases should be much like managing any other system with fluidity and cross-functionality. Working with data should feel consistent to all other parts of the application which should encourage anyone to feel they can engage with and support the database environment. Although it's impossible to prevent 100% of failures, successful DBRE's are able to build the necessary guard rails and educate their developers so that they no longer have a fear of destroying things in case something goes wrong.

BMC Solutions

BladeLogic Database Automation improves efficiency, lowers costs, and reduces risk by automating routine administrative tasks and compliance processes. IT operations can accelerate

deployment and patching, reduce security vulnerabilities, enforce compliance, and achieve greater productivity by streamlining the database management process. Therefore, businesses can stop worrying about the state of their databases and focus on driving growth.

Key Benefits

- Cut database provisioning time by up to 95%
- Reduce manual tasks for highly qualified resources by 85%
- Expand the ratio of administrators to database without impacting productivity or increasing operational costs
- Scale across physical and cloud-based application environments
- Reduce downtime and improve quality of service by applying patches in a rolling fashion

Learn more about [BladeLogic Database Automation](#).

Control-M for Databases allows you to integrate and automate batch workloads for different database systems from a single point of control. Defining new jobs is easy with a user interface that provides interactive object selection, which minimizes manual input and reduces user errors.

Key Benefits

- Reduces Total Cost Ownership (TCO) with native support for multiple databases plus wizards to guide users, eliminating the need to create and maintain scripts and utilities, along with related costs
- Protects database integrity by eliminating complexity and errors through intuitive job forms and job definition validation features, assuring database job accuracy
- Maintains compliance by securing the authentication and connection criteria required for database job submission

Learn more about [Control-M](#) and [Control-M for Databases](#) or [Try Control-M for Free](#).